

METHOD AND APPARATUS FOR IDENTIFYING A MAXIMUM FRAME SIZE

Cross Reference to Related Applications

This application is a continuation of application

5 serial number 08/996,115, entitled "METHOD AND APPARATUS FOR
IDENTIFYING A MAXIMUM FRAME SIZE TO MAINTAIN DELAY AT OR
BELOW AN ACCEPTABLE LEVEL" filed on by David Langley on
December 22, 1997, and is related to the subject matter of
application serial number 08/672,674 entitled, "Autosensing
10 LMI Protocols in Frame Relay Networks" filed on June 28,
1996 by Shankar Natarajan and Gregory Fowler having the same
assignee as this application and incorporated herein by
reference in its entirety.

Inventor

15 David Langley

Express Mail Label Number

EL528759700US

Attorney Docket Number

1211

20 Field of the Invention

The present invention is related to computer networks
and more specifically to network management software in a
computer network.

Background of the Invention

Computer networks such as frame relay or X.25 networks transmit binary information from an originator to one or more recipients. An endpoint device such as a router 5 receives the information from the originator and partitions the information into chunks of information known as frames. During the partitioning process, header information such as routing information is added to the frame based upon instructions from the originator. The router can partition 10 the information into any size frames it desires, subject to a maximum allowable frame size for the network. The endpoint device then queues the frames and transmits them into the network one at a time.

The network uses the routing information to bring the 15 information sent by the originator to a destination endpoint device that accepts information for a specified recipient. The destination endpoint device then reassembles the information from the frames and passes the information to the recipient. Each endpoint device is any device that can 20 supply and/or receive information from the network, such as a conventional router with a frame relay interface or a conventional computer with a conventional frame relay

network interface card installed along with the appropriate software.

Some endpoint devices, such as conventional routers, are shared by more than one originator. A router may receive information to transmit from a second originator before it has transmitted the information from a first originator. Because the frames from the two originators will share the same communication line from the router to the frame relay network, a conflict for the shared resource of the communication line can exist when more than one originator tries to use the endpoint device to send information. If the router can multitask, it can perform the routing and partitioning function while it transmits information from the first originator. However, because the communication line from the router to the network can only handle one frame at a time, the router will append the frames from the second originator into the queue for transmission at a later time, such as after the information from the first originator is transmitted into the network.

Some information is time-sensitive and therefore a simple first-in, first-out queue can cause the time sensitive information to be delayed longer than is acceptable. Such delay can produce undesirable results.

For example, if the information sent to the router by the second originator described above is time sensitive, and the information sent to the router by the first originator is not time sensitive, the time sensitive information must wait 5 until the non-time-sensitive information is fully transmitted.

Examples of time-sensitive information in a network are voice and video data. Voice data is time sensitive because the recipients may notice any delay in the conversation, and 10 video data is time sensitive because, once started, any interruptions will be viewed as either "snow" or "jerky" motions. While some amount of delay is acceptable in both cases, the acceptable limit of the delay is finite and small. The acceptable delay can be exceeded due to a large 15 data transmission behind which the time-sensitive information can be placed in the queue of the endpoint device. The acceptable delay is typically expressed in units of time, such as 50 milliseconds for voice data.

Some frame relay networks allow information to be prioritized in the queue. Time sensitive information 20 received after non-time-sensitive information will nevertheless be framed, and those frames will be moved ahead in the queue of the endpoint device and transmitted into the

network immediately, being delayed only for so long as to allow any partially transmitted frames to complete transmission.

In some circumstances, this delay to wait for partially transmitted frames to complete transmission can still be unacceptable if the time it takes to complete transmission of a partially transmitted frame exceeds the acceptable delay. The time it takes to complete transmission of any partially transmitted frame, referred to herein as transmission delay, is a function of the transmission speed of the endpoint device into the network and the number of remaining bits of the frame in the process of transmission when the time-sensitive information is received. In the worst case, the remaining number of bits is one less than the size of the frame.

One solution to this problem is to use an unshared endpoint device to transmit time-sensitive information. However, such a solution may be prohibitively expensive. Another potential solution is to obtain a very fast endpoint device and connection into the network for which the transmission delay for frames which are even the maximum allowable frame size would be below the acceptable delay. However, this solution may also be prohibitively expensive.

Another potential solution is to immediately abort the transmission of the frame in process if time-sensitive information is received. However, such a solution would cause the recipient to request retransmission of the aborted frame, increasing the traffic on the network, potentially significantly. Another potential solution would be to reduce the frame size on the network so that, in all cases, the transmission delay will be less than the acceptable delay. However, the reduced frame size would mean an increased number of frames. Because each frame contains a certain amount of overhead in the header such as the routing information, the increased number of frames could significantly increase the network traffic.

It is desirable to allow originators of time-sensitive information to transmit frames within the acceptable delay using endpoint devices and network connections of any speed and that are shared with other information without significantly affecting traffic of a frame relay network.

Summary of Invention

A method and apparatus allows a two endpoint devices in communication to negotiate the size of frames to be sent in each direction between an originator and a recipient. Each endpoint device can receive an acceptable delay for each

originating device that uses the endpoint device to send information over the frame relay network. The endpoint device uses the lowest of the acceptable delays and the speed of the line to calculate the maximum size frame it can 5 send and receive in order to maintain delays below the lowest acceptable level. Each endpoint device then transmits the maximum frame size it calculates to other endpoint devices on the frame relay network from which it is capable of receiving frames. Each endpoint device on the 10 network sends frames that are not larger than the lower of the sending endpoint device's maximum frame size and the receiving endpoint device's maximum frame size.

The resulting frames are the largest size that will not exceed an acceptable level of delay for either the sending 15 or receiving device, minimizing the network traffic impact without exceeding the acceptable delay of the sending and receiving endpoint device. The frame sizes are tailored to the needs of each pair of endpoint devices: endpoint devices in communication for which the maximum possible delay is 20 acceptable can use the maximum allowable frame size for the network, such as 4096 bytes. Other endpoint devices in communication will use a lower size frame, but not lower than necessary to ensure a delay not greater than the acceptable delay, minimizing the network traffic impact to

that which is necessary to accommodate the needs of each device in communication.

Brief Description of the Drawings

Figure 1 is a block schematic diagram of a conventional computer system.

Figure 2 is a block schematic diagram of an apparatus for identifying the frame size to be used to transmit information over either of two PVCs to prevent delay from exceeding one or more acceptable levels according to one embodiment of the present invention.

Figure 3 is a flowchart illustrating a method of identifying the frame size to be used to transmit information over each of at least one PVC to prevent delay from exceeding one or more acceptable levels according to one embodiment of the present invention.

Detailed Description of a Preferred Embodiment

The present invention may be implemented as computer software in an endpoint device. An endpoint device may be implemented as a conventional computer system.

Referring now to Figure 1, a conventional computer system 150 for practicing the present invention is shown. Processor 160 retrieves and executes software instructions

stored in storage 162 such as memory, which may be Random Access Memory (RAM) and may control other components to perform the present invention. Storage 162 may be used to store program instructions or data or both. Storage 164, 5 such as a computer disk drive or other nonvolatile storage, may provide storage of data or program instructions. In one embodiment, storage 164 provides longer term storage of instructions and data, with storage 162 providing storage for data or instructions that may only be required for a shorter time than that of storage 164. Input device 166 10 such as a computer keyboard or mouse or both allows user input to the system 150. Output 168, such as a display or printer, allows the system to provide information such as instructions, data or other information to the user of the system 150. Storage input device 170 such as a conventional floppy disk drive or CD-ROM drive accepts via input 172 15 computer program products 174 such as a conventional floppy disk or CD-ROM or other nonvolatile storage media that may be used to transport computer instructions or data to the system 150. Computer program product 174 has encoded thereon computer readable program code devices 176, such as magnetic charges in the case of a floppy disk or optical 20 encodings in the case of a CD-ROM which are encoded as

program instructions, data or both to configure the computer system 150 to operate as described below.

Referring now to Figure 2, one embodiment of an apparatus that provides a frame size to avoid exceeding one or more acceptable delays is shown according to one embodiment of the present invention. Administration 228 receives from input 229 the acceptable delay, measured in units of time such as milliseconds, for each of two PVCs, PVC A and PVC B. The acceptable delay may be calculated to take into account any network delay. The apparatus shown in figure 2 accommodates two PVCs, however any number of PCs may be accommodated in other embodiments of the present invention. Administration 228 stores the acceptable delay received for PVC A in acceptable delay storage 224A and the acceptable delay for PVC-B in acceptable delay storage 224B. Administration 228 receives at input 229 the local speed, in a number of bits per second, such as kilobits per second, of the physical connection 202 used to carry the traffic for the PVCs. This local speed may be the speed of the transceiver 232 that transmits and receives information to and from the frame relay network over communications line 202, the speed of a switch to which the line 202 is connected, or any other speed. Administration 228 stores the local speed received in local speed storage 226. A

single physical connection 202 is shown in Figure 2,
although any number of physical interfaces 202 may be
accommodated by the present invention.

Local maximum size calculator 222 calculates the
5 maximum frame size that may be transmitted or received
without exceeding either of the two acceptable delays stored
in acceptable delay storage 224A, 224B. In one embodiment,
the maximum frame size is calculated by selecting the
minimum of the acceptable delays for all PVCs over which the
10 endpoint device may communicate, and multiplying the
selected delay by the speed of the line to produce a result
having units of a number of bytes.

In one embodiment, local maximum size calculator 222
calculates the maximum frame size by identifying the smaller
15 of the acceptable delays stored in acceptable delay storage
224A, 224B and multiplying the smallest acceptable delay
identified with the speed of the physical connection 202
stored in local speed storage 226. In another embodiment, a
constant value is subtracted from the result of this
20 multiplication to accommodate the delay the endpoint device
requires to calculate the route. The maximum frame size
calculated is stored by local maximum size calculator 222 in
local maximum size storage 218. If a new or different

acceptable delay is received by administration 228, administration 228 signals local maximum size calculator 222 to recalculate and store in local maximum size storage 218 the maximum frame size.

5 The local maximum frame size is also passed by local maximum size calculator to frame size message generator 236.

Frame size message generator 236 receives from administration 228 an identifier of each PVC defined to the

endpoint device 200. These identifiers may be received via input 229 supplied by a person who manages the endpoint

device or received automatically via transceiver 232 over the network. Either periodically or at the time of any

change, frame size message generator 236 generates, for each PVC defined to administration 228, a message indicating the

local maximum frame size, and an identifier of the PVC.

Frame size message generator 236 passes the messages to

protocol adapter 230, which adapts the frame size message

into the protocol required to be transmitted to the other endpoint devices which can receive messages over the PVC

20 specified in each of the messages. In one embodiment,

protocol adapter 230 places the message into a conventional frame relay frame and marks the frame as a frame size

message. The protocol used can be any protocol that allows the receiving device to determine that the message is

management information intended for the endpoint device, such as the protocol described in copending application serial number 08/672,674, entitled "Autosensing LMI Protocols in Frame Relay Networks" filed on 6/28/96 by 5 Shankar Natarajan and Gregory Fowler and having the same assignee as the present invention, which is incorporated by reference herein in its entirety.

Transceiver 232 transmits the frame size message over the PVCs to the frame relay network to be delivered to the other endpoint devices with which the endpoint device 200 communicates. In one embodiment, if no maximum acceptable delay is defined for any PVC, no frame size message is generated and transmitted by the endpoint device 200. In another embodiment, if the frame size in the frame size message would be larger than a certain threshold, such as 1600, 4096 or a maximum allowable frame size for the network, frame size message generator 236 does not produce such frame size messages and therefore no frame size message is generated and transmitted by the endpoint device 200.

20 Transceiver 232 receives frame size messages from other endpoint devices. In one embodiment, the frame sizes received are calculated by different endpoint devices as described above, using acceptable delays and speeds

applicable to each other endpoint device. Transceiver **232** passes such messages to protocol adapter **230**, which decodes them using the protocol described above and passes them to frame size message receiver **234**. Frame size message

5 receiver **234** identifies the frame size contained in each frame size message, and identifies whether such a maximum frame size corresponds to PVC-A or PVC-B using the PVC identifier contained in each frame size message and the PVCs defined to administration **228**. Frame size message receiver **234** stores in distant maximum size storage **220A, 220B** the frame size received for PVC-A and PVC-B, respectively.

In one embodiment, more than one frame size may be received for a PVC. Such embodiment might be useful where many devices share the same PVC. Frame size message receiver **234** compares a frame size received with any frame size stored in the distant maximum size storage **220A, 220B** for that PVC, and stores the frame size received only if it is below the size stored in the corresponding distant maximum size storage **220A, 220B**. In such embodiment, some 20 or all of the frame size messages contain an indicator to allow only the endpoint device that sent the frame size message having the smallest frame size to adjust it upward should conditions warrant.

Minimum selector **216A** selects a minimum of the distant maximum size storage **220A** or the local maximum size storage **218** and stores the result in frame size storage **214A**.

Minimum selector **216B** selects a minimum of the distant maximum size storage **220B** or the local maximum size storage **218** and stores the result in frame size storage **214B**. In one embodiment, a single minimum selector takes the place of the two minimum selectors **216A**, **216B** by selecting the minimum frame size of PVC-A and PVC-B one PVC at a time, and storing the result in the appropriate one of frame size storage **214A**, **214B**.

In one embodiment, at startup of the endpoint device, administration **228** initializes local maximum size storage **218** and distant maximum size storage **220A**, **220B** with the value of the largest possible frame size. The connections to administration **228** are not shown to avoid cluttering the Figure. In this manner, if no acceptable delay is defined or a distant maximum size is not received, it will not impact the size calculation performed by minimum selector **216**.

When information is received for segmenting by conventional frame relay framers **212A**, **212B** via inputs **211A**, **211B**, conventional framers **212A**, **212B** will use the frame

sizes stored in their respective frame size storage **214A**, **214B** as a maximum size to size the frames of the information to be transmitted over PVC-A and PVC-B, respectively.

Framers **212A**, **212B**, such as conventional frame relay

5 assemblers/disassemblers, segment the information received from an originator. Additionally, framers **212A**, **212B** add the conventional frame relay frame header and trailer information and pass the frames to conventional frame relay queue **210**. The functions of framers **212A**, **212B** are described in Smith, "Frame Relay: Principles and Applications" (1993 Addison Wesley, ISBN 0-201-62400-1). In one embodiment, the functions of framers **212A**, **212B** may be combined into a single framer which frames by selecting one of the maximum sizes stored in frame size storage **214A**, **214B** based upon the destination of the information specified by the originator. Frame relay queue **210** queues the frames, using the conventional priority capability described above, and sends them one at a time to transceiver **232** for transmission to the frame relay network over communications line **202** to one or more remote endpoint devices.

20 Referring now to Figure 3, a method of selecting a frame size is shown according to one embodiment of the present invention. Local and distant maximum frame sizes are initialized **310** to the maximum allowable frame size

allowed for the network in one embodiment, and a smaller number in another embodiment. One or more acceptable delays are received and stored 312. In one embodiment, only the lowest value delay is stored in step 312, and in another embodiment, all delays received are stored. A local speed is received and stored 314. The local speed is a speed of communication in one embodiment, which may be equal to a maximum transmission speed of an endpoint device, the maximum speed of transmission of the transmission line coupled to an endpoint device, the maximum speed of reception of a device such as a switch that is coupled to an endpoint device, or any other speed.

The local maximum frame size is calculated 316 as described above by multiplying the delay by the frame size in one embodiment, or multiplying the delay by the frame size and subtracting a constant number from the result. In one embodiment, if more than one acceptable delay is received, the local maximum frame size is calculated as described above using the delay received in step 312 representing the smallest delay.

If the local maximum size calculated in step 316 is above a threshold, it does not need to be transmitted in one embodiment 318. The threshold may be the maximum allowable

frame size for a network in one embodiment, or a different number in another embodiment. If the local maximum size exceeds the threshold, the method continues at step 328, otherwise the method of the present invention continues at 5 step 320.

A PVC is selected 320. In one embodiment, the PVC selected in step 320 is selected from a list of PVCs. The local maximum frame size calculated in step 316 is transmitted 322 to one or more endpoint devices that communicate using the PVC selected in step 320 as described above. If there are additional PVCs 324, for example additional PVCs on the list of PVCs described above, another PVC is selected 326 and the method continues at step 322 until the local maximum frame size is transmitted using some 10 or all the PVCs on the list. 15

One or more distant maximum frame sizes are received and stored 328. In one embodiment, these frame sizes received are calculated using steps similar or identical to 310 - 322, using acceptable delays and speeds applicable to 20 remote devices. In one embodiment, in step 318, the minimum frame size is stored associated with a PVC, and in one such embodiment, an identifier of the PVC with which to associate the distant maximum frame size is received with the frame

size as a part of step 328. The minimum of the frame size received in steps 316, 328 is identified 330, and associated with the PVC with which the distant maximum frame size received in step 328 is associated. The order of steps 318-5 326 relative to steps 328 and 330 need not be as shown in Figure 3.

A payload of information is received 332 and framed 334 using conventional framing techniques described above using the minimum of the maximum frame size identified in step 10 330. In one embodiment, the payload received in step 332 corresponds to a PVC, either implicitly or explicitly using an identifier associated with the PVC. The minimum frame size used to frame the payload in step 334 is the minimum frame size associated with that PVC. The frames of the 15 payload framed in this step 334 are then transmitted 336 over a frame relay network until a new payload is received or a new acceptable delay is received.

If a new acceptable delay is received 338, the method continues at step 312, and if a payload is received 338, the 20 method continues at step 332. In one embodiment, if a new delay is received during the transmission of the frames of a payload, the steps beginning at step 312 are not taken until all the frames of any such payloads are transmitted. In

another embodiment, if a new delay is received during the transmission of the frames of a payload, the steps beginning at step 312 are taken immediately, and steps 332 and 334 are skipped immediately following step 330.